

CLAIMS

I claim:

1. A method of optimizing a response time for retrieving relevant documents from a set of candidate documents identified in response to a search query where the search query includes one or more terms, the method comprising the steps of:

assigning a term weight to each of the terms;

associating a document to a relevance score bin based on a total matched term weight where a document that matches a total term weight of M is associated to a more relevant score bin than a document that matches a total term weight less than M; and

retrieving a set of most relevant documents based on the association to the relevance score bins having a highest relevance score without retrieving other candidate documents.

2. The method as set forth in claim 1 further including determining a relevance of a document independently in relation to other candidate documents based on the relevance score bin associated to the document.

3. The method as set forth in claim 1 further including determining a set of most relevant documents from the candidate documents based on the relevance score bin associated to the documents without determining an exact relevance score for all the candidate documents.

4. The method as set forth in claim 1 wherein assigning the term weight to a term is based on an inverse frequency scoring.
5. The method as set forth in claim 1 further including:
defining a total relevance score range; and
defining one or more relevance score bins within the total relevance score range as a function of a total term weight, the total term weight being a sum of the term weights from the search query.
6. The method as set forth in claim 5 wherein the total relevance score range is divided into the one or more relevance score bins each having an equal size.
7. A computer readable medium having computer executable instructions for performing the method of claim 1.
8. An object retrieval system comprising:
logic for processing a search query having one or more terms;
object retrieval logic for identifying candidate objects that match the search query;
ranking logic for assigning a term weight to each of the terms and associating each combination of matched term weights to a relevance score range, the ranking logic grouping the candidate objects based on the matched term weight where an object that matches a total term weight of M is associated to a more relevant score range than an object that matches a total term weight less than M ; and

retrieval logic for retrieving a set of relevant objects associated to the relevance score ranges having a greatest matched term weight without retrieving the candidate objects from other relevance score ranges.

9. The object retrieval system of claim 8 wherein the logic for processing the search query including a parser that parses the search query to identify the terms.

10. The object retrieval system of claim 8 wherein the ranking logic includes:

logic for defining a total relevance score range; and

logic for defining the relevance score ranges within the total relevance score range as a function of possible term weights that an object can match.

11. The object retrieval system of claim 8 wherein the retrieval logic includes logic for retrieving only objects from the candidate objects that match a highest value of the term weights.

12. The objects retrieval system of claim 8 wherein the ranking logic includes means for associating document relevance scores to the relevance ranges based on matched term weight.

13. A method of retrieving most relevant documents from a set of candidate documents that match a search query having one or more terms, the method comprising:

...the ... of ...

15. The method as set forth in claim 13 wherein retrieving the most relevant documents includes retrieving documents that match a highest term weight associated to a highest score bin without retrieving documents associated to other score bins.

{ KXP0489.DOC;1 }

